Warning Concerning Copyright Restrictions

The Copyright law of the United States (Title 17, United States Code) governs the making of photocopies or other reproductions of copyright material. Under certain conditions specified in the law, libraries and archives are authorized to furnish a photocopy or other reproduction. One of these specified conditions is that the photocopy or reproduction not be "used for any purposes other than private study, scholarship, or research." If a user makes a request for, or later uses, a photocopy or reproduction for purposes in excess of "fair use," that user may be liable for copyright infringement.

CHAPTER 8

The Representation of Information in Long-Term Memory (Part II)

APPLICATION: **More Studying and Daydreaming**

MODELS OF SEMANTIC MEMORY
- *Hierarchical Network Model*
- *Spreading Activation Model*
- *ACT Theory*
- *Feature-Comparison Model*
- *Schema Theory*

THE ROLE OF IMAGERY
- *Facilitative Effects of Interactive Images on Memory*
- *Paivio's Dual-Code Hypothesis*
- *Evidence that Images Are Analog Representations*
- *The Propositional View of Pylyshyn*
- *Kosslyn's Updated Dual-Coding Position*
- *ACT Theory's Treatment of Imagery*
- *Imagery and the Associative-Cognitive Debate*

CONCLUDING REMARKS: **The Maltese Cross Model Revisited**

APPLICATION: More Studying and Daydreaming

Following dinner, Alan returns to his preparation for the upcoming unit exam in his biology course. As Alan reviews the various classes of vertebrates, his consideration of the distinguishing features of the bird class leads him quite involuntarily to begin mentally listing a number of specific examples of birds. Before long, Alan is thinking about ostriches, even though information about ostriches is not a part of the to-be-tested material. He is reminded of the fact that they have long thin legs, they are tall, and they can't fly. He even experiences a mental picture of an ostrich running, flapping its wings, and failing to become airborne. Thinking about long thin legs, in turn, leads Alan to begin thinking about the human female anatomy. When he finally realizes that his mind has wandered, he shifts his focus of study to the anatomy of the human heart, another topic to be covered on the test.

In this case, Alan quite deliberately creates a vivid mental picture of the heart and its surrounding structures in order to promote memory for the information. Later, when Alan reviews the heart again just before going to bed, he first generates a mental picture of the heart as a means of prompting his listing of the various to-be-remembered structures.

These mental experiences of Alan involve that component of long-term memory identified in the preceding chapter as semantic memory. They also involve the use of mental images, which a number of memory theorists believe to be another important component of long-term memory. In this chapter, we shall examine how learning and memory psychologists have attempted to describe the representation of semantic knowledge in long-term store and the role that images play in long-term memory representations.

MODELS OF SEMANTIC MEMORY

Despite the continuing debate over whether semantic and episodic knowledge should be conceptualized as separate and distinct systems within the long-term store that was described in the preceding chapter, a number of formal models of semantic memory have been proposed. These models have in common the aim of describing how decontextualized conceptual knowledge might be represented in long-term store.

Hierarchical Network Model

Collins and Quillian (1969, 1972) proposed one of the first formal models of semantic memory, based in part on the results of experiments using the *speeded verification task.* In this task, sentences tapping subjects' general knowledge (e.g., *A canary is a bird*) are presented, and the amount of time it takes subjects to respond "true" or "false" is measured. The time taken to respond to the sentences is assumed to reflect the organization or structuring of information in memory, because the organization of information will influence how quickly the relationship between concepts can be located in memory and evaluated.

In order to explain the pattern of reaction times found, Collins and Quillian proposed that concepts are arranged hierarchically and interconnected by labeled associations, as illustrated in Figure 8–1. They proposed three types of associations or links: *isa,* representing a subset relation; *has,* representing an attribute relation; and *can,* representing another

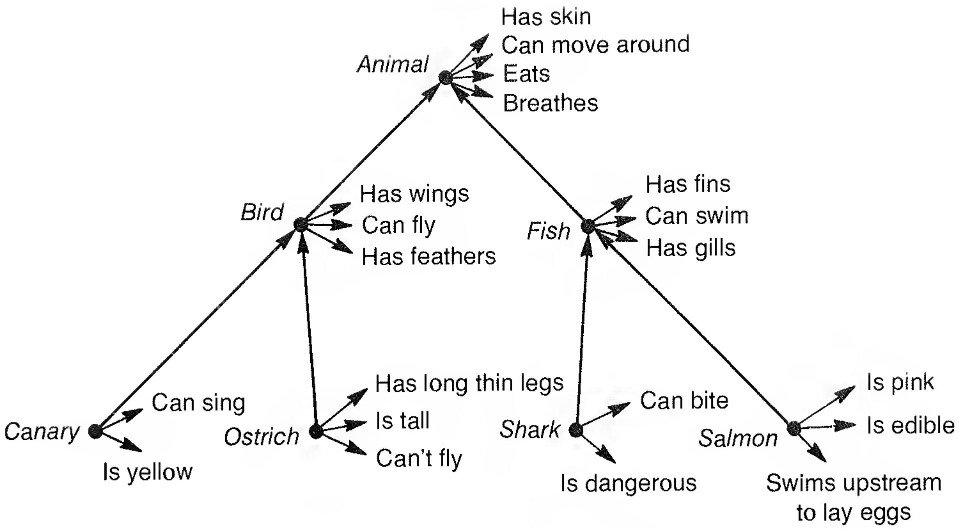

FIGURE 8–1 An example of the hierarchical organization of concepts in semantic memory proposed by Collins and Quillian (1969) in their hierarchical network model of semantic memory.

Source: Collins, A. M., & Quillian, M. R. (1969). Retrieval time from semantic memory. *Journal of Verbal Learning and Verbal Behavior, 8,* 241. Reprinted with permission.

type of attribute relation. Alan's daydreaming about ostriches can be explained within the context of the Collins and Quillian model by assuming that thinking about the concept *birds* led to the activation of the concept *ostrich* by means of an *isa* link. In turn, activating *ostrich* led to the consideration of attributes such as *has long thin legs* by means of the *has* attribute relation.

In proposing a hierarchical arrangement of concepts, Collins and Quillian incorporated into their model what is known as the **cognitive economy assumption.** This assumption stipulates that each property is stored only at the highest possible level in the hierarchy. For example, in Figure 8–1, the properties *has skin* and *eats* are represented in the hierarchy only at the level of *animal* and not at the lower levels of *bird, fish,* or *shark.* Thus responding to ''true'' statements in the speeded verification task is often an inferential process in that a number of associative pathways may have to be traversed in finding an associative link between the concepts.

The search process is based on the assumption that it takes a constant amount of time to move from one level to another in the hierarchy and that the verification time is therefore based on the number of associative pathways that have to be traversed in finding the link between concepts. These assumptions were supported by early experiments of Collins and Quillian showing, for example, that it took subjects longer to verify *A canary has skin* than *A canary is yellow,* and it took subjects longer to verify *A canary is an animal* than *A canary is a bird.*

What are the most important distinguishing features of the hierarchical network model of semantic memory?

Problems with the Hierarchical Network Model

Familiarity/Relatedness Effects. Other evidence, reported by Rips, Shoben, and Smith (1973), quickly challenged some of the assumptions of the hierarchical network model. They found, for example, that subjects are faster in verifying *A bear is an animal* than in verifying *A bear is a mammal,* despite the fact that *bear* and *mammal* are closer in the hierarchy than are *bear* and *animal.* That is, subjects are faster in responding to the more familiar concept *animal* than to the less familiar concept *mammal.* Apparently, subjects perceive *bear* and *animal* to be more closely related than *bear* and *mammal.*

Typicality Effects. Subjects also verify a dominant or salient member of a category consistently more rapidly than they do a less salient item (Wilkins, 1971). For example, *A robin is a bird* is verified faster than *A chicken is a bird.* These findings are inconsistent with the model in that the model predicts equal verification times because the instances of a category are equidistant from the main category node or memory location (see Figure 8–1).

Associative Frequency Effects. Similarly, Conrad (1972) has argued that the strength of preexisting associative links between items being processed in a verification task is a powerful determinant of reaction time. Take, for example, the noun-property pairs, *salmon is edible* and *salmon can move around.* In the Collins and Quillian hierarchy, *is edible* is closer to *salmon* than is *can move around.* However, research suggests that *is edible* is no more strongly associated with *salmon* than is *can move around.*

Conrad (1972) obtained what are called **production frequency measures** by asking subjects to list all of the properties associated with a concept such as *salmon* and by counting the number of subjects who produced each property (hence, the term *production frequency*). When other subjects were asked to verify the two statements, verification times did not differ. Thus Conrad (1972) has argued that the strength of the noun-property association would appear to be more important than the hierarchical distance between the concepts in determining verification times. The fact that *shark eats* is verified as rapidly as *animal eats* can be explained on the same basis. The typicality and familiarity/relatedness effects cited above can also be viewed as associative frequency effects, and all of these findings argue against the cognitive economy assumption of the hierarchical network model.

Why are findings of familiarity/relatedness effects, typicality effects, and associative frequency effects inconsistent with the hierarchical network model?

Acquisition of General Properties. Collins and Loftus (1975) have also argued that the manner in which the general properties of a category are acquired argues against the strong form of cognitive economy incorporated in the model. They note that the knowledge of general properties is based on specific instances involving subordinate items. Thus, during acquisition, specific properties are likely to be multiply represented. Even as properties common to most instances of a category are generalized to the category as a whole, they may, in some instances (as suggested by evidence cited above), also remain strongly associated with specific instances of the category.

Responding to False Sentences. Evidence that certain false statements are rejected more slowly than others is also a problem for the model. For example, *A St. Bernard is a cat* is rejected relatively slowly despite the absence of hierarchical associative pathways between the units. Thus the slowness of response would appear to be a function of the semantic relatedness of the concepts. The model must be increased in complexity to cope with such negative instances.

How is the hierarchical network model deficient in explaining the acquisition of general properties and the response to false sentences?

Spreading Activation Model

In order to deal with the problems associated with the hierarchical network model, Collins and Loftus (1975) proposed a revision of the model. The revised model retained the associative network but rejected the strict hierarchical structure of the original model. An illustration of the revised model's representation of semantic knowledge is shown in Figure 8–2.

Organization of Concepts

The strict hierarchical structure of the original Collins and Quillian model was replaced by a structure based on the concept of **semantic distance** (or **semantic relatedness**) in order to account for familiarity/relatedness effects. In this model, lengths of associative links are used to represent the strength of the relationship between concepts, with strength determi-

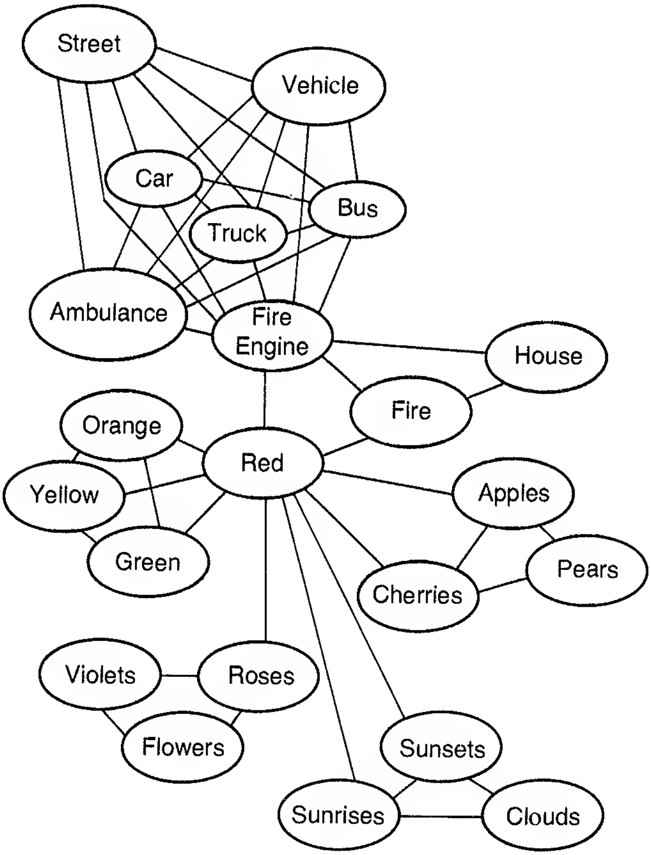

FIGURE 8–2 An example of the interrelationship of concepts in semantic memory as represented in Collins and Loftus's (1975) spreading activation model of semantic memory. In this representation, the length of the associative links between the concepts represents the strength of the relationship.

Source: Collins, A. M., & Loftus, E. F. (1975). A spreading-activation theory of semantic processing. *Psychological Review, 82,* 412. Copyright 1975 by the American Psychological Association. Reprinted by permission.

nations based on production frequency norms. Other methods for determining relatedness include rating word pairs for relatedness and rating typicality of category instances. On this model, in representing the semantic distance between *bird* and *ostrich* and between *bird* and *robin* in Alan's semantic memory store, the associative link between *bird* and *ostrich* would probably be longer.

The model's rejection of a strict hierarchical arrangement of concepts also involved a rejection of the cognitive economy assumption. The Collins and Loftus model allows for much more prestored knowledge such that less inferencing based on searches of the network is required. As an example of prestored knowledge, the model also includes the additional associative link of *isnota.* Such a link, it is argued, is needed to account for fast false responses to items such as *A bat is a bird.* The speed of the response suggests that the response is not based on the search of the network, but instead reflects direct accessing of the unit *A bat is* not *a bird.*

How does the spreading activation model differ from the hierarchical network model in its description of how information is organized in semantic memory?

Retrieval Processes in the Spreading Activation Model

This spreading activation model of Collins and Loftus assumes a more complicated retrieval process than does the hierarchical network model. Activation of a given concept is assumed to spread to related concepts. Thus the activation of *bird* in Alan's semantic memory store spreads to related concepts such as *ostrich,* which in turn spreads to related concepts such as *has thin long legs,* which in turn spreads to *human female anatomy.* This last example of spread of activation within Alan's memory can be accounted for, in part, because the spreading activation model does not assume a strict hiearchical arrangement of concepts.

This spread of activation notion replaces the original model's simple retrieval process of moving from one node to another. As the activation spreads, its strength decreases. Thus concepts farther from the originally activated node are less likely to be activated than are closer concepts. The spread of activation is also dependent on the strength of initial activation. Greater initial activation of a concept results in greater spread of activation from that concept. The model also assumes that links between concepts differ in accessibility or strength, based, for example, on their frequency of use, and that the strength of a link determines the speed with which activation spreads.

These assumptions thus provide a basis for explaining typicality effects. *A robin is a bird* is verified faster than *A chicken is a bird* because the more frequent accessing of the *robin-bird* associative link has made that link stronger, more accessible, and more quickly traversed at the time of retrieval. In a verification situation, the system is monitoring for the intersection of activation spreading from the concepts contained in the statement, because an intersection is evidence that the activated concepts are related in some way. If an intersection of activation is detected, decision processes evaluate the total strength of the activation involved in the intersection. If the total activation exceeds a criterion or threshold value, the type of intersection found is evaluated and a response is executed.

Collins and Loftus (1975) also incorporated into the model a provision for mutually exclusive links, in accordance with suggestions by Holyoak and Glass (1975). (See Glass & Holyoak [1975] for a description of their related marker-search model of semantic memory.) Such an assumption helps explain why subjects are so fast at rejecting as false statements such as *A mallard is an eagle.* Presumably *mallard* and *eagle* share the superordinate *bird.* Thus the model assumes that when activation spreads from *mallard* and *eagle* to *bird,* the activation of links that are mutually exclusive provides strong evidence in favor of responding "false."

The model also incorporates a counterexample strategy, again based on the work of Holyoak and Glass (1975), as a source of negative evidence in evaluating statements. Subjects' very fast rejection of statements such as *All birds are canaries* can be explained by assuming that, once the superordinate relationship between *canaries* and *birds* is confirmed, an attempt is made to generate another subordinate of *birds* that is mutually exclusive from *canaries.*

✎ *How does the spreading activation model explain retrieval from semantic memory?*

Priming and Semantic Context Effects

The model provides an explanation for semantic priming effects in lexical decision tasks. A **lexical decision task** can involve a subject's deciding whether a group of letters is a word. For example, a subject might see in succession the following letter groups:

nwat
red
truck
taleb

The subject's task would be to respond to each letter group as quickly as possible with "yes" or "no," indicating that the letter group is or is not a word. If the speed with which a subject responds "yes" to *truck* in the above sequence is compared to the speed with which a subject responds "yes" to *truck* in the following sequence, a clear difference is usually found:

nwat
car
truck
taleb

The response to *truck* in the second sequence is usually faster (e.g., see Huttenlocher & Kubieck, 1983; Meyer & Schvaneveldt, 1971; Neely, 1977).

This effect whereby the processing of one word facilitates the processing of another word related to it in meaning is an example of **semantic priming.** The notion of a spread of activation to related concepts is consistent with this finding. The spread of activation from *car* to *truck* would mean that *truck* would already be somewhat "turned on." Thus when the word *truck* itself is presented, its representation in memory already stands out somewhat. Thus, the memory representation of *truck* will be accessed and matched against the incoming letter pattern faster than if *car* had not been presented previously.

Furthermore, the notion of semantic relatedness or semantic distance can account for the fact that preceding *truck* with *car* produces a larger priming effect than does preceding *truck* with *vehicle.* As represented in Figure 8–2, the length of the association between *truck* and *car* is shorter than the length of the association between *truck* and *vehicle,* meaning that there will be a greater transfer of activation between *car* and *truck.*

This basic notion of spread of activation can also help account for semantic context effects in general. *As Mary accelerated onto the freeway, the car in front of her suddenly stopped. Reflexively, Mary stepped on the _____. In the process of reading these sentences, the individual might activate her or his schema for driving, which would result in a spread of activation to concepts, such as brake,* contained within that schema. Thus the reader might only have to analyze superficially the sensory input associated with the next word in the sentence in order to confirm the expectation that *brake* did occur (Lindsay & Norman, 1977; Rumelhart & Siple, 1974).

How can the spreading activation model account for priming and semantic context effects?

Assessment of the Model

As was indicated in the preceding description, the model's flexible assumptions regarding both the representation and processing of information enable it to account for many of the phenomena that could not be explained easily by the original hierarchical network model. However, the model's flexibility is a weakness as well as a strength. Kintsch (1980) has essentially argued that the model is so general as to make it impossible to derive critical predictions that can be tested. In the absence of such critical tests, the adequacy of the model is difficult to assess.

Likewise, Chang (1986) has noted that while the model can explain phenomena such as typicality effects and the slowness of responding to false statements containing related concepts it does not strongly predict such effects. In addition, Johnson-Laird, Herrmann, and Chaffin (1984) have criticized the spreading activation model, as well as other models of semantic memory, for its failure to account for various important components of semantic processing. For example, the model, they claim, does not explicitly account for relationships such as synonymy and antonymy. However, Chang (1986) does emphasize the potential utility of this type of general explanatory approach in observing that the model may serve as an overall framework within which more specific models of semantic representation and processing may be incorporated.

What are the major strengths and weaknesses of the spreading activation model?

ACT Theory

In considering the Adaptive Control of Thought (ACT) representation of semantic memory, it must be remembered that Anderson has consistently rejected the assumption of distinct episodic and semantic long-term stores (Anderson & Bower, 1973; Anderson, 1983; Anderson, 1990). Therefore, the representation of semantic information described in the following section would also apply, according to ACT, to episodic information. As was noted in the preceding chapter in the discussion of ACT's distinction between procedural and declarative knowledge, ACT assumes that declarative knowledge is represented in the form of a network of propositions, as illustrated in Figure 7–1 for episodic declarative information. As is shown in Figure 8–3, this same sort of propositional network can be used to represent semantic information.

Critical to ACT's description of semantic memory, as it is to Collins and Loftus's model, is the concept of spread of activation. Once a unit in the propositional network is activated, the activation will spread along associative pathways to associated concepts. The concepts most affected by this spread will be those most closely linked with the original site

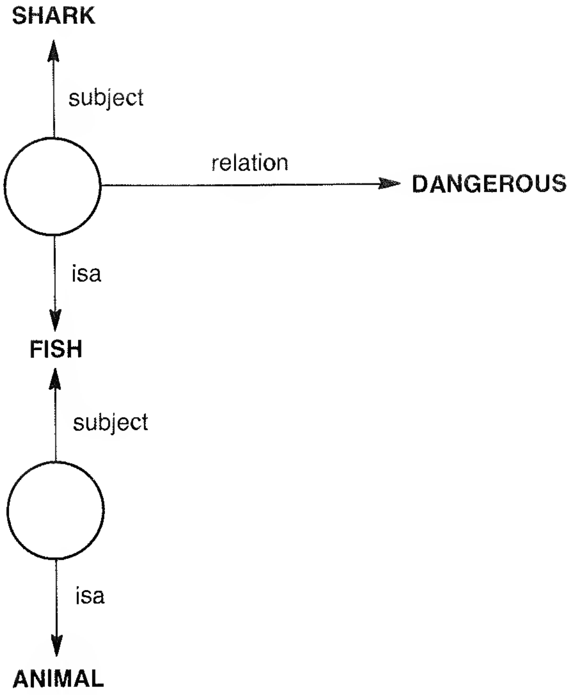

FIGURE 8–3 Semantic knowledge represented as a network of propositions.

of activation. Thus, in these general terms, ACT's emphasis on spread of activation is almost identical to Collins and Loftus's spreading activation model.

☑ *According to ACT, how are declarative episodic and semantic information represented in long-term memory?*

Spread of Activation and the Execution of Productions

ACT goes beyond the spreading activation model in attempting to describe more precisely what determines the closeness of links between concepts within the network and thus what determines the spread of activation. In his update of ACT, Anderson (1983) argues that each concept node (where node refers to a memory location) has an associated strength based primarily on its frequency of use.

ACT also assumes that the spread of activation affects the execution of productions. Recall that productions are condition-action pairs. The action is performed only when certain conditions, as represented in declarative memory, are met. A pattern matcher, consisting of a sequence of pattern tests, monitors for a match between activated nodes in declarative memory and the conditions of the production. ACT stipulates that the degree of activation of a particular pattern-matcher node is a function of the strength of that node (based on its frequency of application), the level of activation of the declarative memory structures being matched, and the degree of match between the memory structures and the conditions embodied in the pattern matcher.

As activation spreads through the declarative memory network, its spread affects the application of productions linked to particular declarative memory structures. For example, Table 8–1 shows productions that Anderson (1983) has incorporated into ACT to help explain the effects of semantic priming on lexical decision tasks. In the case of automatic lexical decisions (i.e., where subjects are judging whether sequences of letters are words but are not consciously anticipating the occurrence of a particular word), if the word to be judged (say, for example, *doctor*) has been preceded by a related word (e.g., *nurse*), the spread of activation in the declarative memory network from *nurse* to related items will speed the rate at which productions 2 and 3 can be applied. If, on the other hand, *doctor* is preceded by an unrelated word, such as *rocket*, the execution of productions 2 and 3 will be unaffected.

What major factor determines the spread of activation through ACT's propositional network, and how does the spread of activation affect the execution of productions?

Explanation of Semantic Verification Results

Anderson has also used the notions of spread of activation and execution of productions to account for semantic verification results of the type described in the presentation of the hierarchical network model and the Collins and Loftus spreading activation model. Anderson (1983) argues that the recognition of a fact associated with a particular concept will take

TABLE 8–1 Example Productions for Performance in the Lexical Decision Task Based on Anderson's ACT System

1. A word-naming production

P1 IF the word is spelled F-E-A-T-H-E-R
THEN assert that the word is similar to FEATHER.

2. Productions that perform automatic lexical decision

P2 IF the goal is to judge if the stimulus is spelled correctly
and a word is similar to the stimulus
and the stimulus mismatches the spelling of the word
THEN say no.

P3 IF the goal is to judge if the stimulus is spelled correctly
and a word is similar to the stimulus
and the stimulus does not mismatch the spelling of the word
THEN say yes.

3. Productions that perform deliberate lexical decision

P4 IF the goal is to judge if the stimulus matches an anticipated word
and a word is anticipated
and the stimulus does not mismatch the spelling of the word
THEN say yes.

P5 IF the goal is to judge if the stimulus matches an anticipated word
and a word is anticipated
and the stimulus mismatches the spelling of the word
THEN change the goal to judging if the stimulus is correctly spelled.

4. An optional production that capitalizes on nonwords similar to the anticipated word

P6 IF the goal is to judge if the stimulus matches an anticipated word
and a word is anticipated
and the stimulus is similar to the word
and the stimulus mismatches the spelling of the word
THEN say no.

Source: Reprinted by permission of the publishers from *The Architecture of Cognition* by John R. Anderson, Cambridge, MA: Harvard University Press, Copyright © 1983 by the President and Fellows of Harvard College.

longer the more facts there are associated with the concepts. This argument is based on the assumption that the more facts there are associated with a concept, the less activation each of the pathways leading from the concept will receive as a result of the activation of the concept itself. On this basis, ACT can explain why Collins and Quillian found that *A canary is a bird* is verified faster than *A canary is an animal.* Presumably, there are in memory more associative links leading from *animal* than from *bird,* and, as a result, the execution of the productions involved in matching the test sentence against declarative memory structures will be slower.

ACT can also explain the typicality and semantic relatedness effects that were such a problem for the hierarchical network model in terms of the strength of nodes and the resulting spread of activation. *A bear is an animal* is verified faster than *A bear is a mammal,* because the strength of the *animal* node, and thus the degree of activation of links leading from it, is greater than the strength of the *mammal* node. In this case, it has to be assumed that the greater strength of the *animal* node more than compensates for the greater number of links leading from it. Otherwise, based on the argument given for the *canary* example, one would have to predict that *A bear is a mammal* would be verified faster.

Anderson (1983, 1984) explains the slow rejection of false statements such *A St. Bernard is a cat* in terms of coincidental interconnections between the concepts. The larger the number of associative interconnections between the concepts, the longer it will take to evaluate all of the connections because of the smaller amount of activation each connection will receive. Anderson suggests a waiting model to explain the rejection of such false statements. That is, the model does not assume that all interconnections are evaluated. Rather, if the execution of productions has not resulted in verification of the statement within a particular time period, a "false" response is made.

✎ *How does ACT account for the original Collins and Quillian findings, typicality/semantic relatedness effects, and the slow rejection of false statements?*

Assessment of the Model

Clearly, ACT provides a framework for explaining many of the effects found in research on semantic memory. Furthermore, it must be remembered that ACT is intended as a framework for cognitive processes in general and is in no way limited to the semantic memory paradigm. In fact, in the final major section of this chapter, ACT will again be discussed in conjunction with imagery effects, and it will be considered in some detail again in Chapter 9 on retrieval from long-term memory. The scope of ACT as an explanatory framework is both its strength and its weakness. Anderson's systematic program of research has shown that ACT can be adapted to account for a wide variety of cognitive activities. However, as was mentioned earlier, ACT has become increasingly complex as an explanatory framework. As a result, it is difficult to determine what sorts of findings would count as evidence against the model.

Relevant to this issue is what Smith (1978) calls the **sufficiency/transparency tradeoff.** Models such as ACT and the spreading activation model of Collins and Loftus strive for sufficiency in their explanation of semantic memory in that they incorporate more and more detailed assumptions about representation and processing of information in order to explain more and more discrete findings. The problem with this trend, Smith notes, is that the addition of details to the model makes it less transparent. That is, the proliferation of details in the model makes it harder and harder to discern the major assumptions and claims of the model. In turn, the harder it becomes to identify the major predictions of the model, the harder it becomes to know what sort of evidence will tend to disconfirm the model. Smith argues that theorists proposing complex theories of cognitive functioning have a responsibility to identify clearly the major assumptions of the model and what would constitute evidence contrary to the model.

What is the major problem with comprehensive explanations of cognitive processes such as ACT?

Feature-Comparison Model

The Set-Theoretic Approach

The feature-comparison model of Smith and colleagues (Smith, Shoben, & Rips, 1974; Smith, 1978) represents an alternative to the network theories of semantic memory. It is an example of set-theoretic models in which concepts are represented by sets of features or attributes. (See Meyer [1970] for a description of his predicate-intersections set-theoretic model.) In the set-theoretic approach, the contents of semantic memory are the attributes of the objects represented. Concepts do not exist independent of the features of which they are composed. The representation of *bird,* for example, might include features such as *is living, can sing, can fly, has feathers,* and so on. Thus there is an emphasis on computed rather than prestored knowledge. In verifying *A robin is a bird,* for example, the model assumes that the relationship is computed at the time of verification by comparing the lists of features associated with the two concepts. Features can be viewed as component parts of an object or values on a dimension. Dimensions represented might include perceptual properties (e.g., shape, size, color), functional characteristics (e.g., mode of locomotion, eating habits), or even abstractions (e.g., honesty, beauty).

This model's emphasis on computed rather than prestored knowledge means that the feature-comparison model is less associative and more cognitive in nature than is the spreading activation model. Whereas the spreading activation model emphasizes the associative links that exist in long-term memory, the feature-comparison model emphasizes mental processes that operate on the stored representation of concepts.

Central to the model is the distinction between defining and characteristic features. **Defining features** are those central to the meaning of a concept. To be considered a member of a particular category, the object must have the defining features. All objects in a particular class share the defining features for that class. *Has feathers* and *is living,* for example, might be defining features of *bird.* **Characteristic features,** on the other hand, are attributes typical of the category but not necessary to its definition. *Can fly* and *can sing* are characteristic of *bird* in that most birds exhibit these properties, but certainly not all birds can fly or sing.

How does the set-theoretic approach to semantic memory differ from network theories, and how does the feature comparison model distinguish between defining and characteristic features?

The Comparison Process

Based on this distinction between defining and characteristic features, the feature-comparison model postulates a two-stage process of feature comparison for the situation in which an individual is asked to verify a statement, such as *A canary is a bird.* Stage 1 involves matching all of the features, regardless of whether they are defining or characteristic, for the relevant concepts. If a large number of features match, a quick "true" response is made. If very few features match, a quick "false" response is made. Within the system, decision processes set the criteria for what constitutes high and low feature similarity. Thus the criteria can vary from situation to situation. When an intermediate number of features are shared, stage 2 of the comparison process must be executed. In stage 2, only the defining features are matched. If all defining features match, a true response is made. If any defining features do not match, a false response is made. As stage 2 requires additional comparison time, decision times will be slow when feature similarity is neither very high nor very low and stage 2 has to be executed.

How does the feature-comparison model's comparison process operate?

Application of the Model

On the basis of this rather elegantly simple two-stage process, the model can account for typicality and familiarity/relatedness effects. According to the model, characteristic features play a crucial role in the effects. For example, consider verification of the statements *A robin is a bird* and *A chicken is a bird.* As members of the category *bird, robin* and *chicken* have the same defining features. Thus, differences in verification time must be due to differences in characteristic features. *Robin* has more features characteristic of the *bird* category and thus can be verified rapidly based on high feature similarity following the stage I comparison. *Chicken,* on the other hand, has fewer features characteristic of the *bird* category and thus the feature similarity found in stage 1 is less likely to exceed the criterion for making a "true" response. As a result, verification would be slowed because of the need to execute stage 2.

Smith and colleagues (1974) also argue that there is evidence from the study of linguistic hedges (e.g., *a true, technically speaking,* and *loosely speaking*) that is consistent with their model. The hedge *a true* is used when the subject noun contains both defining and characteristic features of the predicate noun, as in *A robin is a true bird. Technically speaking* is used when the subject noun contains the predicate noun's defining features but not its characteristic features, as in *Technically speaking, a chicken is a bird. Loosely speaking* is used when the subject and predicate nouns do not have matching defining features but do share a number of characteristic features, as in *Loosely speaking, a bat is a bird.* In the case of negative statements in which the subject and predicate nouns share many characteristic features but do not share defining features, *technically speaking* is used, as in *Technically speaking, a bat is not a bird.*

Smith (I978) argues that the general set-theoretic approach in which concepts are represented by sets of features or attributes also provides a framework for explaining some effects of context on word meanings. For example, Barclay, Bransford, Franks, McCarrell, and Nitsch (1974) had subjects process the word *piano* in the context of one of two sentences: *The man lifted the piano* or *The man tuned the piano.* Then, subjects received cues for recall of all of the sentences processed. Barclay and colleagues found that *heavy* was a better cue than *musical* for the first sentence, whereas the reverse was true for the second sentence. Smith suggests that the word *lifted* in the first sentence would increase the saliency of the feature *weight* in the representation of *piano* and thus the saliency of the *weight* feature in the composite semantic representation of the entire sentence. In the second sentence, the feature of *music-producing* in the representation of *piano* and in the composite representation of the sentence would be enhanced by the word *tuned.* Thus the semantic features of contextual information may influence subtly the feature representation of particular words without in any way fundamentally changing their meaning.

In another study of context effects on meaning, Anderson and Ortony (1975) found that subjects interpreted the word *container* to mean *basket* when it occurred in the sentence *The container held the apples* but interpreted it to mean *bottle* when it occurred in the sentence *The container held the cola.* In the case of the second sentence, Smith (1978) argues that characteristic features of *cola* are added to the features of *container* and result in the activation of the more specific concept *bottle.* In the case of the first sentence, Smith proposes that the function of *container* is determined by examination of its features and that features of *apples* relevant to this function are then added to the representation of *container,* resulting in the identification of the more specific concept *basket.*

What semantic memory effects are well explained by the feature-comparison model?

Problems with the Model

Despite the model's success is accounting for many phenomena of semantic verification and natural language use, a number of criticisms have been leveled against it. One central challenge involves the issue of defining features. Various researchers (e.g., McCloskey & Glucksberg, 1978; Rosch & Mervis, 1975) now argue that natural categories, in fact, may not have defining features. That is, there may be no features that are common to all members of the category. Instead, Rosch and Mervis (1975) suggest that natural semantic categories be conceptualized as networks of overlapping attributes. The greater the extent to which a category member bears a family resemblance to other members of the category, the more prototypical it is judged to be.

When McCloskey and Glucksberg (1978) asked subjects to decide whether items belonged to particular categories, they found that subjects disagreed with each other in judging intermediate-typicality items, as in deciding, for example, whether *bookends* is part of the category *furniture,* and that subjects were inconsistent in their own judgments from one session to the next. As a result of these findings, McCloskey and Glucksberg suggest that natural categories are fuzzy sets with no clear boundaries separating members from nonmembers. Thus, to the extent that the concept of defining features is questionable, the processing details of the feature-comparison model require reexamination.

Evidence produced by Holyoak and Glass (1975) in the study of responses to false statements also presents a problem for the feature-comparison model. They found instances in which the usual familiarity/semantic relatedness effect was reversed. For example, they found that *All fruits are vegetables* was rejected more quickly than *All fruits are flowers,* and *Some chairs are tables* was rejected more quickly than *Some chairs are beds.* This finding is contrary to the feature-comparison model because the second instance in each pair would be more likely to require stage 1 processing only. This prediction follows from the stipulation that stage 1 processing terminates in a quick ''false'' response when the number of features in common to the two concepts is so low as to exceed the criterion for low feature similarity. For both example pairs, such low feature similarity would be more likely for the second statement.

Collins and Loftus (1975) argue, however, that this result can be explained by the spreading activation model's assumption that *fruits/vegetables* and *chairs/tables* are labeled as mutually exclusive subordinates, whereas *fruits/flowers* and *chairs/beds* are not. Discovery of such negative evidence would lead to a fast ''false'' response.

In an analysis of other false statements, Holyoak and Glass (1975) also found that statements such as *All animals are birds* are also rejected quickly. Again, this finding is troublesome for the feature-comparison model because the high feature similarity of *ani-*

mals and *birds* would mean that the statement could only be correctly rejected as "false" following the stage 2 comparison. Using production frequency measures, Holyoak and Glass (1975) found that the speed of rejecting such statements was better predicted by people's tendency to produce, in response to the subject noun, a concept that is a counterexample to the predicate noun (e.g., *fish*) than by people's tendency to produce the predicate noun itself. That is, they found that the semantic relatedness of the subject and predicate nouns was less important than the tendency to produce a counterexample to the predicate noun.

Collins and Loftus (1975) cite both of these findings from the research of Holyoak and Glass as evidence of the need to incorporate into models of semantic memory multiple strategies for disconfirming, such as searching for mutually exclusive links and counterexamples. The lack of alternative processing strategies, they contend, is thus one of the limitations of the feature-comparison model.

Yet another set of findings involving semantic relatedness is also problematical for the feature-comparison model. In their research, McCloskey and Glucksberg (1979) attempted to manipulate the criterion for relatedness that might be used in making a stage 1 comparison by presenting false statements in which the subject and predicate nouns were semantically related (e.g., *All birds are canaries*) interspersed in a list also containing true statements. The reasoning here is that subjects experiencing such a list would have to set a very strict criterion for feature similarity in stage 1 in order to avoid making quick erroneous responses to the false statements containing related nouns.

However, if the feature-comparison model is correct, setting such a strict criterion should mean that the usual semantic relatedness effect for true statements should be reduced, because the probability of having to execute stage 2 would be increased. That is, verifying *All robins are birds* should take just about as long as verifying *All penguins are birds* because, even in the case of the *robins* statement, the degree of feature similarity might not be sufficient to exceed the criterion for a "true" response in stage 1. Contrary to this prediction derived from the model, McCloskey and Glucksberg found that the semantic relatedness effect was even more pronounced than in a situation in which the false statements contained unrelated nouns.

In interpreting this evidence against the feature-comparison model, it should be noted that the accumulation of such evidence is to a great extent a function of the model's specificity. Unlike the spreading activation model, the feature-comparison model does clearly generate testable predictions and thus can be falsified. However, the model's specificity is also a function of its limited scope, which certainly makes it far from a comprehensive approach to describing semantic memory. As Cohen (1977) has observed, any network model can represent a variety of relationships and thus can be extended as needed, while it is not clear how set-theoretic models such as the feature-comparison model can represent relationships other than class membership and properties.

Certain aspects of the feature-comparison model may profitably be incorporated into more general network models that allow for greater flexibility in representation and processing. It is entirely possible that the analysis of concepts into component features may be important under certain conditions, but the ecological validity of any model that assumes that feature analysis of concepts is invariant must be questioned (e.g., see Kintsch, 1974). Surely concepts such as *bird* have a unitary prestored representation in the semantic system that under many conditions does not require decomposition into constituent features.

What are major weaknesses of the feature-comparison model?

Property-Comparison Model

In conjunction with their presentation of evidence inconsistent with the feature-comparison model, McCloskey and Glucksberg (1979) proposed an alternative that has become known as the property-comparison model. In this model, McCloskey and Glucksberg subscribe to the set-theoretic approach of describing concepts as sets of properties, but they dispense with the distinction between defining and characteristic features. Each property is composed of attributes, with only typical values of each attribute stored. The model provides for a comparison process in which the subject noun is matched against the predicate noun. As the comparison of attributes proceeds, positive and negative evidence is accumulated, and the probability of the statement's truth or falsity is assessed. Once the probability exceeds some threshold value established for the task at hand, a response is made.

McCloskey and Glucksberg have demonstrated that this model can account for their finding that including in a list false statements in which the subject and predicate nouns are semantically related (e.g., *All birds are canaries*) increased the semantic relatedness effect for true statements. The model can explain this effect by assuming that there is a higher probability that a given comparison will yield positive evidence when, in fact, the statement is false in the situation in which the list contains false statements with semantically related elements. This means that each comparison for true statements that yields positive evidence will contribute less to reaching the "yes" decision threshold.

In turn, more comparisons will have to be made in evaluating true statements in order to reach the decision threshold. This increase in the number of comparisons needed will be greater for true sentences containing weakly related nouns than for true sentences containing strongly related nouns, because the proportion of comparisons that yield positive evidence is lower for true statements containing weakly related nouns. However, the model is hardly comprehensive. For example, Chang (1986) has pointed out that the model does not provide a framework for explaining the verification of single-property statements such as *A bird has skin.*

How does the property-comparison model differ from the feature-comparison model?

Schema Theory

In his commentary on the limitations of models of semantic memory, Kintsch (1980) suggested that investigators must begin to consider units of language beyond the individual word and he proposed schema theory as a profitable framework for extending semantic memory research. Consider again the example schemas shown in Figure 7–4. The representations shown there clearly bear a resemblance to the associative network models that have been considered in the discussion of semantic memory. Recall that Rumelhart and Ortony (1977) assume that schemas are embedded within each other and thus range in specificity from the relatively elemental units shown in Figure 7–4 to very much more generalized knowledge structures such as that for *washing clothes*. Within this approach, concepts are represented in terms of the interrelationships that exist among the constituent elements of the concept (Rumelhart, 1980).

Note also that a schema approach such as that of Rumelhart and Ortony provides a mechanism for explaining context effects of the type discussed in conjunction with the feature-comparison model. The different interpretations of *broke* in the sentences *David* *broke the glass* and *David broke the chain* can be explained in terms of other stored information in memory concerning glasses and chains. As Kintsch (1980) observes:

> *Words obtain their meaning from the scripts, categories, and schemata in which they participate. It is the relatively stable structure of semantic memory that provides a basis for the semantics of words. These larger units ought to be the main concern of semantic memory research. (p. 617)*

It is important to recognize that the schema approach is not incompatible with some of the models of semantic memory previously discussed. The notion of schemas could certainly be incorporated into Collins and Loftus's spreading activation model. It must be remembered that the spreading activation model was proposed to account for a relatively limited body of findings from the semantic memory literature and did not describe the structure of the underlying memory representation with as much specificity as Rumelhart and Ortony have. However, the model could certainly be elaborated to include the more detailed type of propositional networks employed by Rumelhart and Ortony in representing the relationships among concepts. Most importantly, the concepts of spread of activation and schemas are perfectly compatible.

The notion of schemas can also be incorporated within ACT. Although Anderson (1983) has criticized specific schema theories for failing to distinguish between declarative and procedural knowledge, ACT's organization of concepts into networks of propositions in declarative memory is certainly consistent with the types of knowledge structures proposed by Rumelhart and Ortony. In fact, in his textbook of cognitive psychology, Anderson (1990) specifically talks about propositions being organized into schemas.

✎ *How is the notion of schema compatible with many of the associative network models of semantic memory?*

THE ROLE OF IMAGERY

The role of imagery in the representation of information in long-term memory has been extensively debated by cognitive psychologists in the last 20 years. The crux of the debate is whether it is necessary to assume that there exist in memory imagery representations of information that are separate and distinct from propositional representations.

Facilitative Effects of Interactive Images on Memory

Although the role of images in the representation of information in long-term memory is a continuing source of debate, the facilitative effect on memory of encouraging subjects to employ visual imagery is well established. For example, in a series of experiments in 1972, Bower demonstrated the beneficial effects of imagery instructions on subjects' paired-associates (P-A) learning. You will recall that, in P-A learning, the subject has to try to recall the second member of each pair in response to the presentation of the first member of the pair.

In one experiment, Bower demonstrated that subjects receiving instructions encouraging them to image the two members of each pair interacting in some way recalled significantly more items than subjects receiving no special study instructions. In another experiment, imagery subjects also significantly outperformed a group of subjects explicitly instructed to repeat or rehearse each presented pair. This finding suggested that instructions to image do not have their effect simply by encouraging verbal processing of the items.

In a third experiment in the series, Bower investigated the influence of a simultaneous motor task on the use of imagery in learning P-As. While trying to learn a list of P-As using imagery, half of the subjects performed a visual tracking task in which they had to follow with two fingers an irregular, wavy black line being fed through a machine at a constant rate. While simultaneously trying to learn the P-As using imagery, the other half of the subjects performed a tactile tracking task in which they had to follow with two fingers a string glued to the paper in an irregular, wavy pattern. Bower reasoned that if instructions to image do indeed result in visual processing the visual tracking task should interfere more with the learning of the P-As than should the tactile tracking task. Consistent with this prediction, he found that the tactile tracking group recalled more P-As than did the visual tracking group.

Finally, Bower investigated the importance of the nature of the image formed. One group received instructions to image the members of the pair interacting, while the other group received instructions to form separate images of each member of each pair. Only the interactive imagery group outperformed the subjects from the first experiment who received no special study instructions. Clearly, these results suggest that creating interacting visual images can facilitate memory for verbal materials.

In the example of Alan, presented at the beginning of the chapter, his attempt to remember the relationships involving the heart and surrounding structures was based on the generation of interactive visual images. In Chapter 10, further evidence of the facilitative effect of imagery on memory will be presented, but first a consideration of theoretical explanations of the effect is in order.

✎ *How did Bower's research on P-A learning point to the important role of imagery in facilitating memory for verbal materials?*

Paivio's Dual-Code Hypothesis

In order to explain the beneficial effects of imagery, Paivio (1971, 1986) has proposed the dual-code hypothesis. This explanation argues that information may be stored in long-term memory in two forms: verbal codes and image codes. The verbal code is available for any event or object that can be described. The image code is available for any event or object that can be visualized. Thus some items can be remembered through only one of the codes, whereas others can be remembered through both. For example, concrete nouns such as *houseboat* and *table* can be represented through both codes, whereas abstract nouns such as *democracy* and *truth* can be represented only through the verbal code. Thus, Alan was able to use imagery, as well as verbal representations, to represent the heart and its surrounding structures because they can be visualized, but he might have had difficulty using imagery to try to represent the biochemical process responsible for the contraction of the heart.

Information represented by two codes, Paivio argues, has a retrieval advantage. Representing information with more than one code increases the probability of locating and retrieving the information at the time of test. Multiple retrieval routes to the information are established. Thus imagery instructions of the type employed in Bower's experiments increase the probability that the information will be multiply represented in long-term memory and therefore more likely to be successfully retrieved. Central to the dual-code hypothesis, of course, is the assumption that imagery representations exist in memory that are separate and distinct from verbal (or propositional) representations.

What are the major assumptions of Paivio's dual-code hypothesis?

Evidence that Images Are Analog Representations

In investigating the assumption that there are distinct imagery representations, researchers have attempted to demonstrate that images have *analog properties.* That is, they have attempted to demonstrate that imagery representations contain visual information that is processed in a manner similar to the processing of visual perceptual information.

Shape Comparisons

Shepard and Chipman (1970) asked subjects to judge the shape similarity of states. One group was given only the names of the states to be judged. The other group was given picture outlines of the states. The two groups were found to use similar criteria in judging the states in that they separated the states into similar groupings. Furthermore, the name-only group reported using imagery in making the judgments. These results suggest that imaging and visually perceiving (as in actually looking at the shape outlines of the states) are similar activities.

Image Scanning Effects

Image scanning effects have been interpreted as supporting the same conclusion. Kosslyn, Ball, and Reiser (1978) asked subjects to hold an image of a map in memory and to imagine a black dot moving from one location to another on the map. Subjects were instructed to push a button when the dot arrived at the second location. Kosslyn and colleagues found that time to push the button was greater when the actual distance between the locations on the map was greater. This finding suggests that spatial information in images of maps may be represented and processed much as it would be represented and processed in visually inspecting the maps.

How are the findings from the study of subjects judging the shape similarity of states and subjects scanning maps consistent with the view that imaging and visually perceiving are similar activities?

Selective Interference Effects

Other evidence of the analog properties of images comes from studies of selective interference effects. In one such study, Brooks (1968) used a block letter such as that shown in Figure 8–4. Subjects were asked to form an image of the letter, and then while holding the image in memory, subjects had to scan the image and decide whether each corner was a part of the very top or the very bottom of the figure. Some subjects were asked to respond verbally, and others were asked to respond by pointing at Ys and Ns arranged in an irregular pattern, as shown in Figure 8–4.

Brooks found that responding visually took twice as long as responding verbally. In order to rule out the possibility that responding visually takes longer regardless of what else the subject is doing, Brooks also tested subjects in a situation in which they had to perform a verbal task and respond either verbally or visually. For that task, responding verbally took

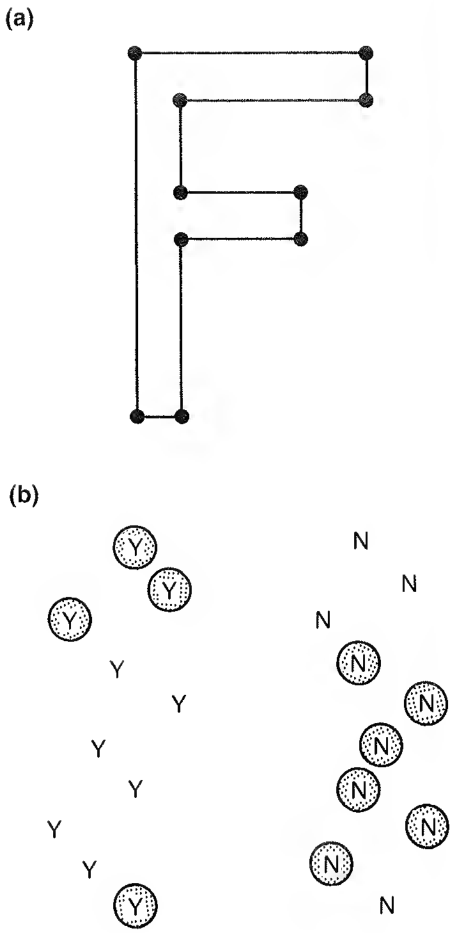

FIGURE 8–4 (a) The type of block letter shown to subjects in the image scanning study by Brooks (1968). (b) The irregular pattern of Ys and Ns to which subjects had to point in the visual response condition of the Brooks study. The shaded circles indicate the correct pattern of responses for the letter *F*.

Source: Brooks, L. R. (1968). Spatial and verbal components in the act of recall. *Canadian Journal of Psychology, 22*, 350–351. Reprinted with permission.

longer. Thus Brooks's results suggest that imaging and perceiving visually (as required in the visual-pointing task) are similar activities in that they interfere with each other.

What evidence did Brooks (1968) find that imaging interferes with visual perception?

Mental Rotation Effects

Mental rotation effects reported by Shepard and Metzler (1971) can also be interpreted as supporting the view that imagery representations contain visual information that is processed in a manner similar to the processing of visual perceptual information. In one of their experiments, subjects were timed as they decided if a pair of drawings showed the same three-dimensional objects. Figure 8–5 shows the type of drawings used in the study. In those cases in which the drawings did depict the same object, the difference in orientation of the two drawings was varied from 0 to 180 degrees. Shepard and Metzler found that the greater the difference in orientation of the members of a pair, the longer it took subjects to respond. They interpreted these results to suggest that subjects made the shape comparisons by mentally rotating one object at a constant rate, just as one would rotate and judge an actual physical object.

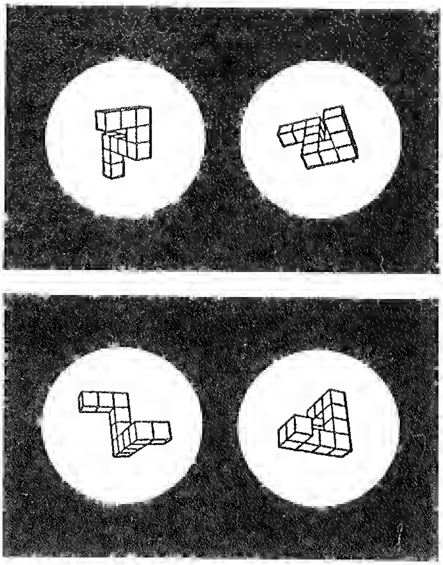

FIGURE 8–5 An example of two pairs of drawings of three-dimensional objects used in Shepard and Metzler's (1971) study of mental rotation effects. The subject's task was to decide if the left- and right-hand member of each pair depicted the same object.

Source: Shepard, R. N., & Metzler, J. (1971). Mental rotation of three-dimensional objects. *Science, 171,* 702. Copyright 1971 by the American Association for the Advancement of Science.

Image Size Effects

A final piece of evidence in support of images having analog properties comes from Kosslyn's (1975) study of image size effects. Subjects were asked to generate either a small or large image of an animal by imaging it in relation to another animal. For example, imaging *dog* in relation to *whale* should result in a relatively small image of dog, whereas imaging *dog* in relation to *wasp* would result in a relatively large image of dog. Then, subjects were asked to judge the appropriateness of various features for the animals imaged (e.g., *Is a beak an appropriate feature of a dog?*). Kosslyn found that the time to judge the appropriateness of features was greater for small images than for large images. This suggests that subjects' performance was influenced by the perceptibility of parts of the image, just as performance might be influenced by the perceptibility of parts of an actual photograph or visual scene.

What results from the study of mental rotation effects and image size effects support the view that imaging and visually perceiving are similar activities?

Collectively, these results suggest that images do contain visual information that is processed in a manner similar to the perceptual processing of visual information. Although such evidence has been widely interpreted as supporting the argument for distinct imagery representations in long-term memory, this view is not universally held, as the next section will show.

The Propositional View of Pylyshyn

While admitting that images may be generated and used in performing specific tasks, Pylyshyn (1973, 1981) nevertheless argues that the basic unit of representation in long-term memory is the proposition. In Pylyshyn's view, all images are based on underlying abstract propositions.

Pylyshyn uses a number of arguments in defending this position. First, he notes that images differ from pictures and physical objects in terms of inaccuracies and forgotten "bits." If something from an image is forgotten, one does not experience a gap in the way one does if a section is torn from a photograph. Instead, the gaps are filled in. For example, if Alan was to generate an image of the area outside the biology building that included the bike area, his image might not contain all of the objects in that area. However, if the fire hydrant next to the bike rack were omitted, Alan would not experience a hole in his image at that location. Instead, the area containing the fire hydrant would be filled in.

In addition, Pylyshyn notes that whole bits tend to be forgotten and that one often knows that bits are missing. Alan is very unlikely to generate an image that contains half of the fire hydrant, but he may be aware that his filled-in image that lacks the fire hydrant is somehow incomplete. Such introspective evidence is consistent, Pylyshyn claims, with the position that an image is based on an abstract underlying propositional representation that includes information about the lost bits.

Pylyshyn also argues that the ability to generate images of events never witnessed before is evidence that the formation of images is dependent on general world knowledge represented in propositional form. Alan, for example, can create an image of an ostrich walking a tight-rope. He is able to do this because he has available general propositional knowledge about giraffes and walking a tight-rope that can be combined into a reasonable image representation.

The facilitative effects of imagery on memory are due, Pylyshyn suggests, to the elaborative propositional processing necessary for image formation. That is, the superior memory resulting from the use of imagery is not due to storing away in long-term memory visual copies of the images formed. Rather, the superior memory is the result of having elaboratively processed the underlying propositional information in order to generate the images. Generating an image of a giraffe walking a tight-rope will necessitate a more extensive consideration of the features and associations connected with *giraffe* and *walking a tight-rope* than will simply repeating the phrase *giraffe walking a tight-rope.* Generating the image will thus make the information more accessible at the time of retrieval.

Some indirect experimental evidence in support of Pylyshyn's position that visual copies of images are not stored in long-term memory comes from a study of picture memory by Nelson, Metzler, and Reed (1974). They compared memory for information contained in a simple sentence (e.g., *A smiling old man holds a little girl*) with memory for the same information presented in the form of a nondetailed drawing, a detailed drawing, and an actual photograph. The results showed the phrase to be more poorly recognized than the three types of pictorial stimuli, but recognition of the three types of pictorial stimuli did not differ. This suggests that the superior memory for pictures was not due to the storing of detailed visual copies in long-term memory. Rather, pictures of any sort result in a more elaborative propositional representation of the information. Pylyshyn makes the same argument for images.

What general arguments has Pylyshyn advanced in arguing against images as a basic form of long-term memory representation?

Kosslyn's Updated Dual-Coding Position

In defending the dual-coding position, Kosslyn (1980) has attempted to identify more precisely the representations and processes that underlie imagery effects. He accepts

Pylyshyn's suggestion that images are sometimes constructed from underlying propositional knowledge. However, he also argues that there are nonpropositional long-term imagery representations. Furthermore, he maintains, these imagery representations have analog properties. Thus Kosslyn argues for both propositional representations and analog codes.

As analog codes, images possess **emergent properties.** That is, they possess properties not easily discerned by a consideration of propositional representations alone. For example, in judging which of two types of bird is prettier, Alan might find images useful because prettiness is an emergent property that could not be easily determined by considering only the propositional information he has stored about the two types of bird (unless, of course, Alan has made this comparison in the past and has recorded his conclusion as a proposition in memory). A series of comparisons of propositional information pertaining to color of plumage, facial features, size, and so on might not permit a quick and easy decision about a holistic property such as prettiness in the way in which a comparison of images might. Thus, in arguing that images have emergent properties Kosslyn stresses the functional role of images in cognitive processing. Not only do images permit comparison of emergent properties but they also permit such valuable cognitive operations as mental rotation and scanning.

It must be stressed here, however, that in arguing for imagery representations, Kosslyn is not suggesting that images are more or less exact mental copies of the objects or events that they represent. Rather, images are in Kosslyn's view **quasi-pictorial,** meaning that they preserve many of the analog properties of pictures but are also more abstract in nature. The most important property preserved is the spatial arrangement of the elements contained in the image. In arguing for the abstract quasi-pictorial nature of images, Kosslyn accepts Pylyshyn's observations that images differ from pictures in important ways.

With regard to the type of information represented in images, a further observation is in order here regarding Brooks's (1968) investigation of selective interference effects in the use of imagery. In addition to the experiment described earlier involving the scanning of the letter and verbal versus visual responding, Brooks also conducted a letter-scanning study in which verbal and spatial responding were compared. For the spatial responding condition, subjects had their eyes closed and scanned with their fingers a series of raised Ys and Ns arranged in the type of irregular pattern shown in Figure 8–4. Subjects were slower to respond in the spatial responding condition than in the verbal responding condition. Thus the interference found in the visual responding condition in the earlier experiment may have been spatial rather than visual in nature. Such a conclusion is consistent with the position that images are abstract spatial representations rather than mental pictures of visual scenes.

It should also be noted that Kosslyn's dual-coding position goes well beyond that of Paivio. In particular, Kosslyn's model (Kosslyn, 1980, 1987) attempts to describe in detail the structures in long-term memory that enable one to use imagery and the cognitive processes that are involved when one uses imagery. For example, Kosslyn's model (Kosslyn, 1980, 1987; Kosslyn, Reiser, Farah, & Fliegel, 1983; Kosslyn, Holtzman, Farah, & Gazzaniga, 1985) includes a description of the type of information stored in long-term memory from which an image in the **visual buffer** (or short-term store) can be generated, using what Kosslyn calls the **picture-processing module.** Such a long-term memory representation might include information about shape (such as lines and angles) in a form that is separate and distinct from the propositional representation of the image.

In addition, Kosslyn's model proposes processes that operate on the underlying image representations. For example, the "put" procedure operates on an image residing in the visual buffer by integrating into the pattern at the correct location the stored image representation of some relevant part. An example is Alan's putting a major blood vessel into an already generated image of the heart or Alan's putting additional tail feathers into an already generated image of a running ostrich. In order for the ''put'' procedure to operate, however, other processing procedures, the **shape-encoding** and **location-encoding processing subsystems,** must first operate to locate the relevant location within the image.

Other research generated by the model suggests that multipart images of three-dimensional objects are generated one part at a time, beginning with nearer surfaces and proceeding to more and more distant surfaces (Kosslyn, Cave, Provost, & von Gierke, 1988; Roth & Kosslyn, 1988). Thus Kosslyn's approach takes the basic dual-coding notion proposed by Paivio and attempts to construct a very specific model of how images are represented and processed within the system.

How does Kosslyn's dual-coding position differ from the propositional view of Pylyshyn and the dual-coding approach of Paivio?

ACT Theory's Treatment of Imagery

John Anderson's position regarding imagery has undergone an interesting evolution. Earlier (e.g., Anderson & Bower, 1973; Anderson, 1976), he argued for propositional representations only. Part of his argument for this position was based on his contention that it is not possible to distinguish between imagery and propositional representations in memory (Anderson, 1976, 1978, 1979). Thus, in pursuit of parsimony, only propositional representations should be assumed.

In more recent versions of ACT (Anderson, 1983), however, Anderson argues for image representations that are separate and distinct from propositional representations. In fact, ACT now embodies three types of representation: abstract propositions, spatial images, and temporal strings (which represent the order of sets of items). Like Kosslyn, Anderson (1983) argues that images most importantly contain information about the orientation and spatial relationships of elements being represented. Thus ACT can explain Pylyshyn's (1973) observation that subjects may be able with certainty to report that the presence of a lamp in their image of a room scene and yet be unable to describe the lamp in detail. The image may represent the relative location of various objects in the room without containing detailed information about the visual properties of the objects.

Just as productions operate on propositional information represented in declarative memory in the ACT model, so too do productions operate on information contained in the image representations. For example, ACT attempts to account for the Shepard and Metzler (1971) mental rotation effects in terms of a production system that operates on the image representations of the two objects being compared.

How does ACT now deal with imagery in describing the representation of information in long-term memory?

Imagery and the Associative-Cognitive Debate

The treatment of imagery in ACT and in Kosslyn's model is yet another example of how the study of learning and memory processes has moved away from a strict associative approach. The representation and processing of images described by the two models is not strictly associative in nature. These models are proposing that the storage and use of images involves much more than the linkage of items in memory. However, the propositional systems assumed by both models also allow for an important contribution of associationist principles. Although Kosslyn does not provide a detailed description of his propositional system, such systems by their very nature tend to be associative, as was seen in many of the descriptions of semantic memory.

How have recent explanations of imagery affected the associative-cognitive debate?

CONCLUDING REMARKS: The Maltese Cross Model Revisited

As noted in the concluding section of the preceding chapter, it is possible to relate many of the issues raised in considering the representation of information in long-term memory to Broadbent's Maltese cross model of memory presented in Chapter 6. For example, it would certainly be possible to incorporate notions such as semantic distance, spread of activation, mutually exclusive associative links, node strength, and imagery representations into the model's treatment of the long-term associative store, even though Broadbent has not yet done so. It would even be possible to incorporate feature comparison processes, property comparison processes, and image manipulating processes and productions into the Maltese cross model's description of the processing system. Before this integration occurs, however, greater consensus is needed as to which of these various notions pertaining to semantic memory and imagery representations are empirically worthy of inclusion in such a general model of memory.